

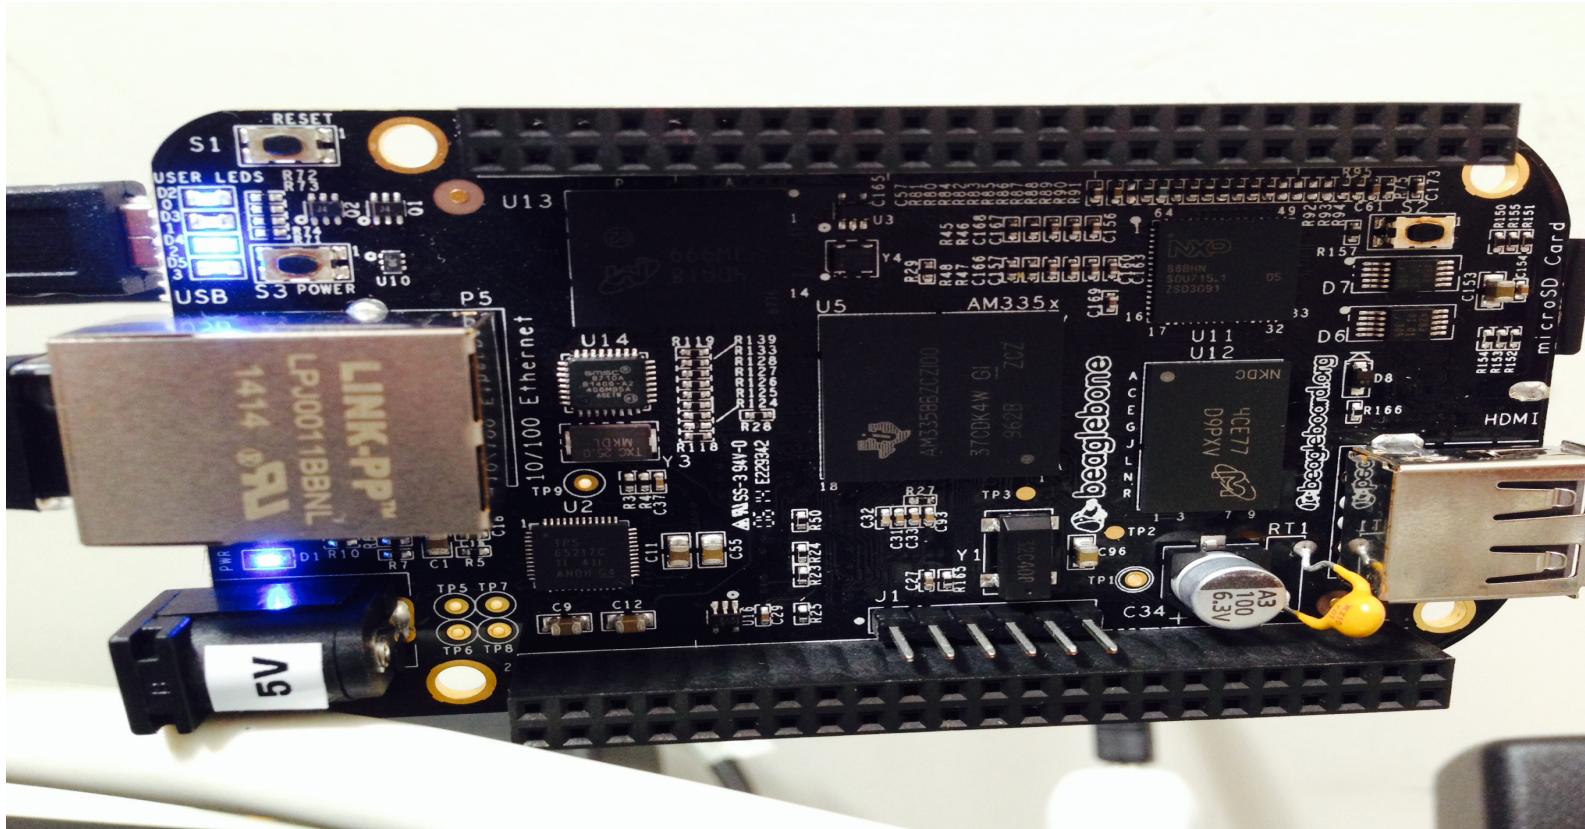
NGINX Powers 12 Billion Transactions per day at Capital One

Rohit Joshi

Mike Dunn



We built a prototype of an API Gateway using a BeagleBone Black, NGINX and Lua

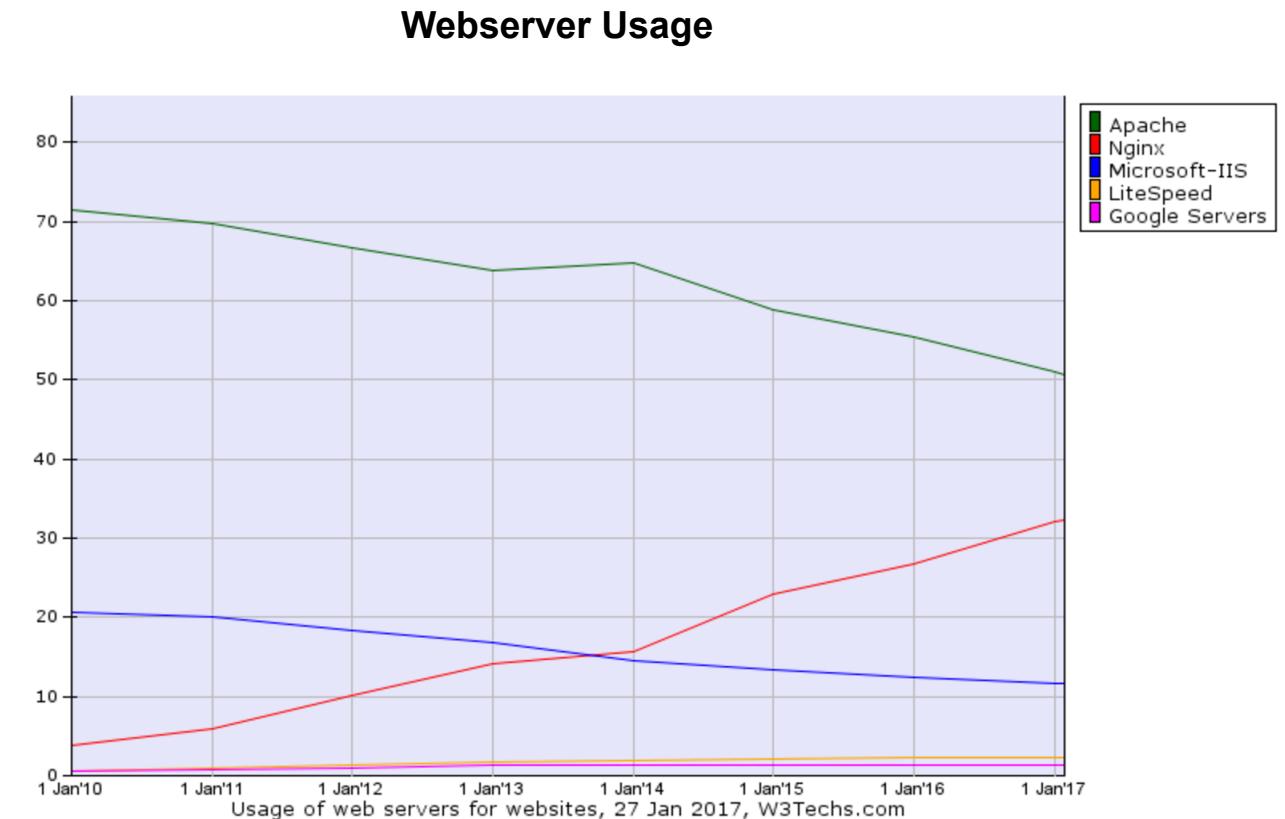


- Beagle Bone Black (5v power)
- ARM Processor
- The prototype was able to sustain:
 - 625 TPS for HTTP
 - 335 TPS for HTTPS
- Nginx PR: [Support for mutual SSL authentication for upstream https proxy](#)

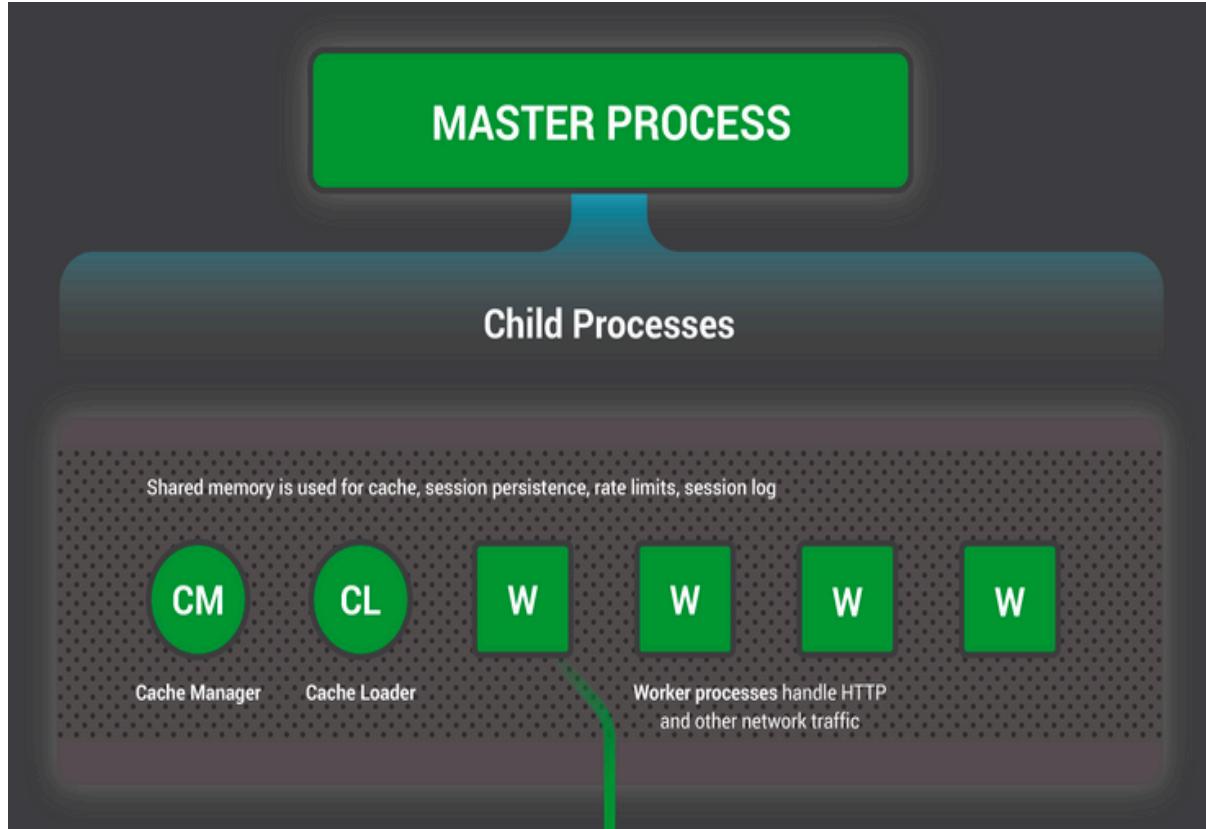
What Technologies did we select for the build and why?

Why NGINX?

- Light Weight : ~10 Mb memory foot print, low CPU Usage
- Concurrent : Supports 100K+ connections
- High performance web server written in C
- Async IO
- Event Driven
- Pluggable architecture
- Native binding to Lua
- Architectural details



Why NGINX?



- The *master* process performs the privileged operations such as reading configuration and binding to ports, and then creates a small number of child processes (the next three types).
- The *cache loader* process runs at startup to load the disk-based cache into memory, and then exits.
- The *cache manager* process runs periodically and prunes entries from the disk caches to keep them within the configured sizes.
- The *worker* processes do all of the work! They handle network connections, read and write content to disk, and communicate with upstream servers.

How NGINX works

NGINX uses a Non-Blocking "Event-Driven" architecture

Listen Sockets & Connection Sockets



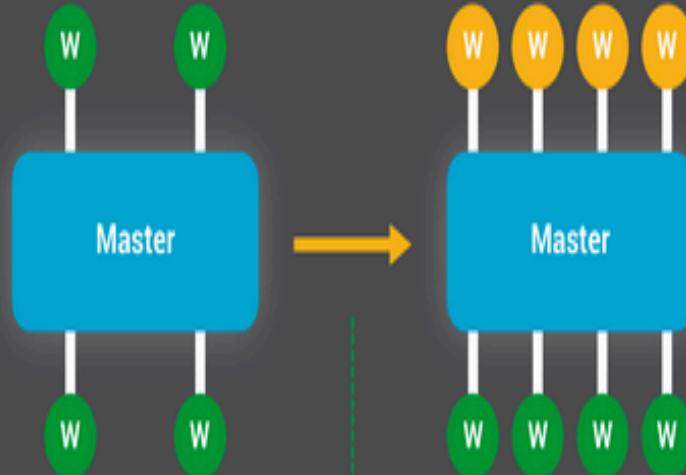
Wait for an event (epoll or kqueue)

- Event on Listen Socket:
- accept new connection
- set connection to be non-blocking
- add connection to the socket list

- Event on Connection Socket:
- data in read buffer? read
- space in write buffer? write
- error or timeout? close & remove from socket list

An NGINX worker can process hundreds of thousands of active connections at the same time

Master starts new worker processes with new configurations

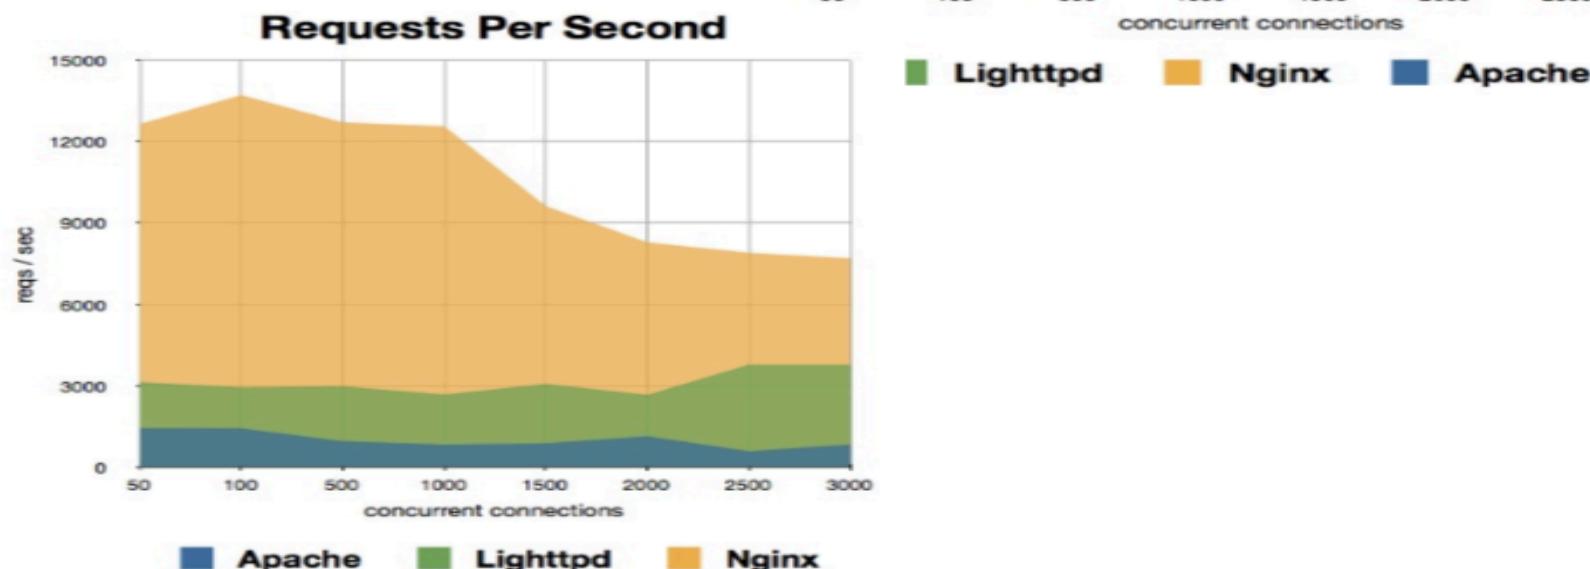
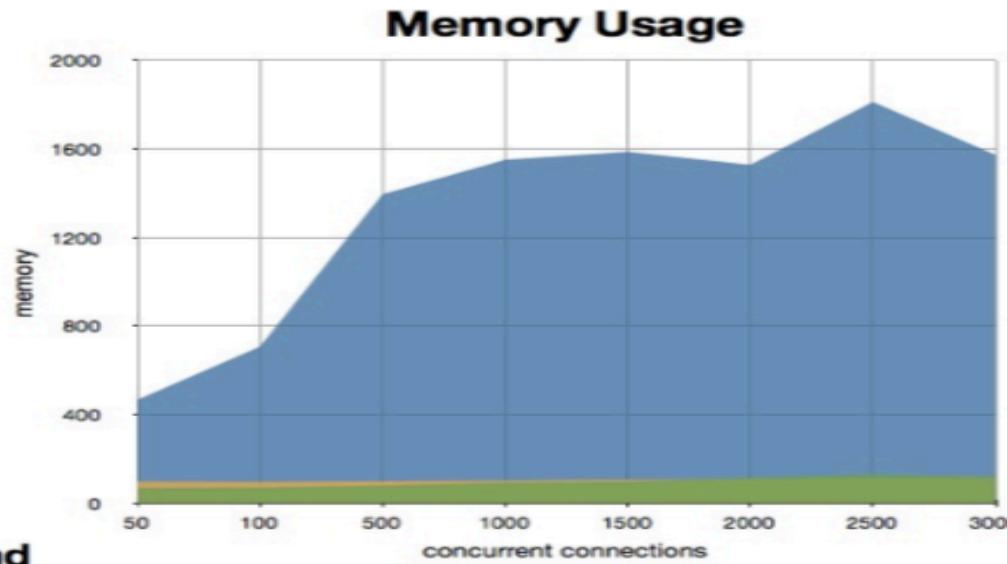


NGINX keeps on running with new configuration, and no interruption in service

Update configuration on disk
SIGHUP master or nginx -s reload

Old worker processes complete existing transactions and then exit gracefully

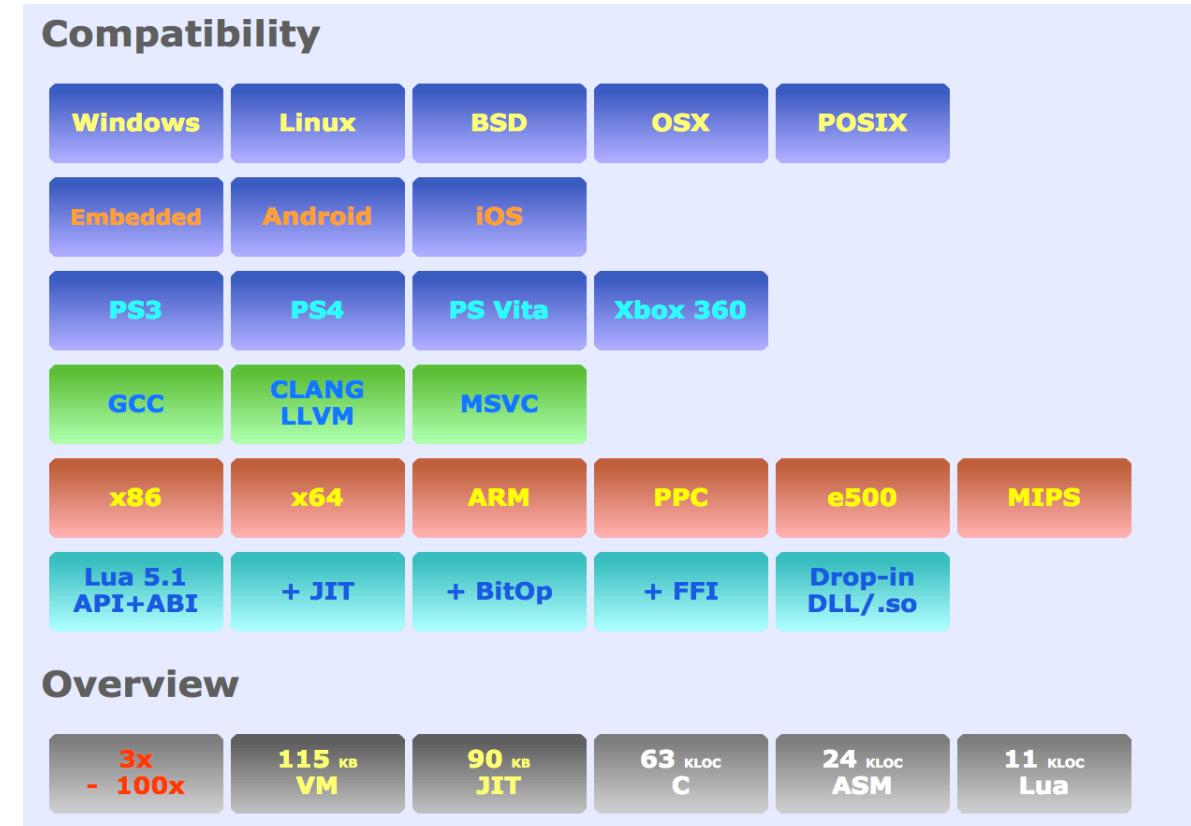
NGINX is fantastic at Scaling and Handling Concurrent Connections



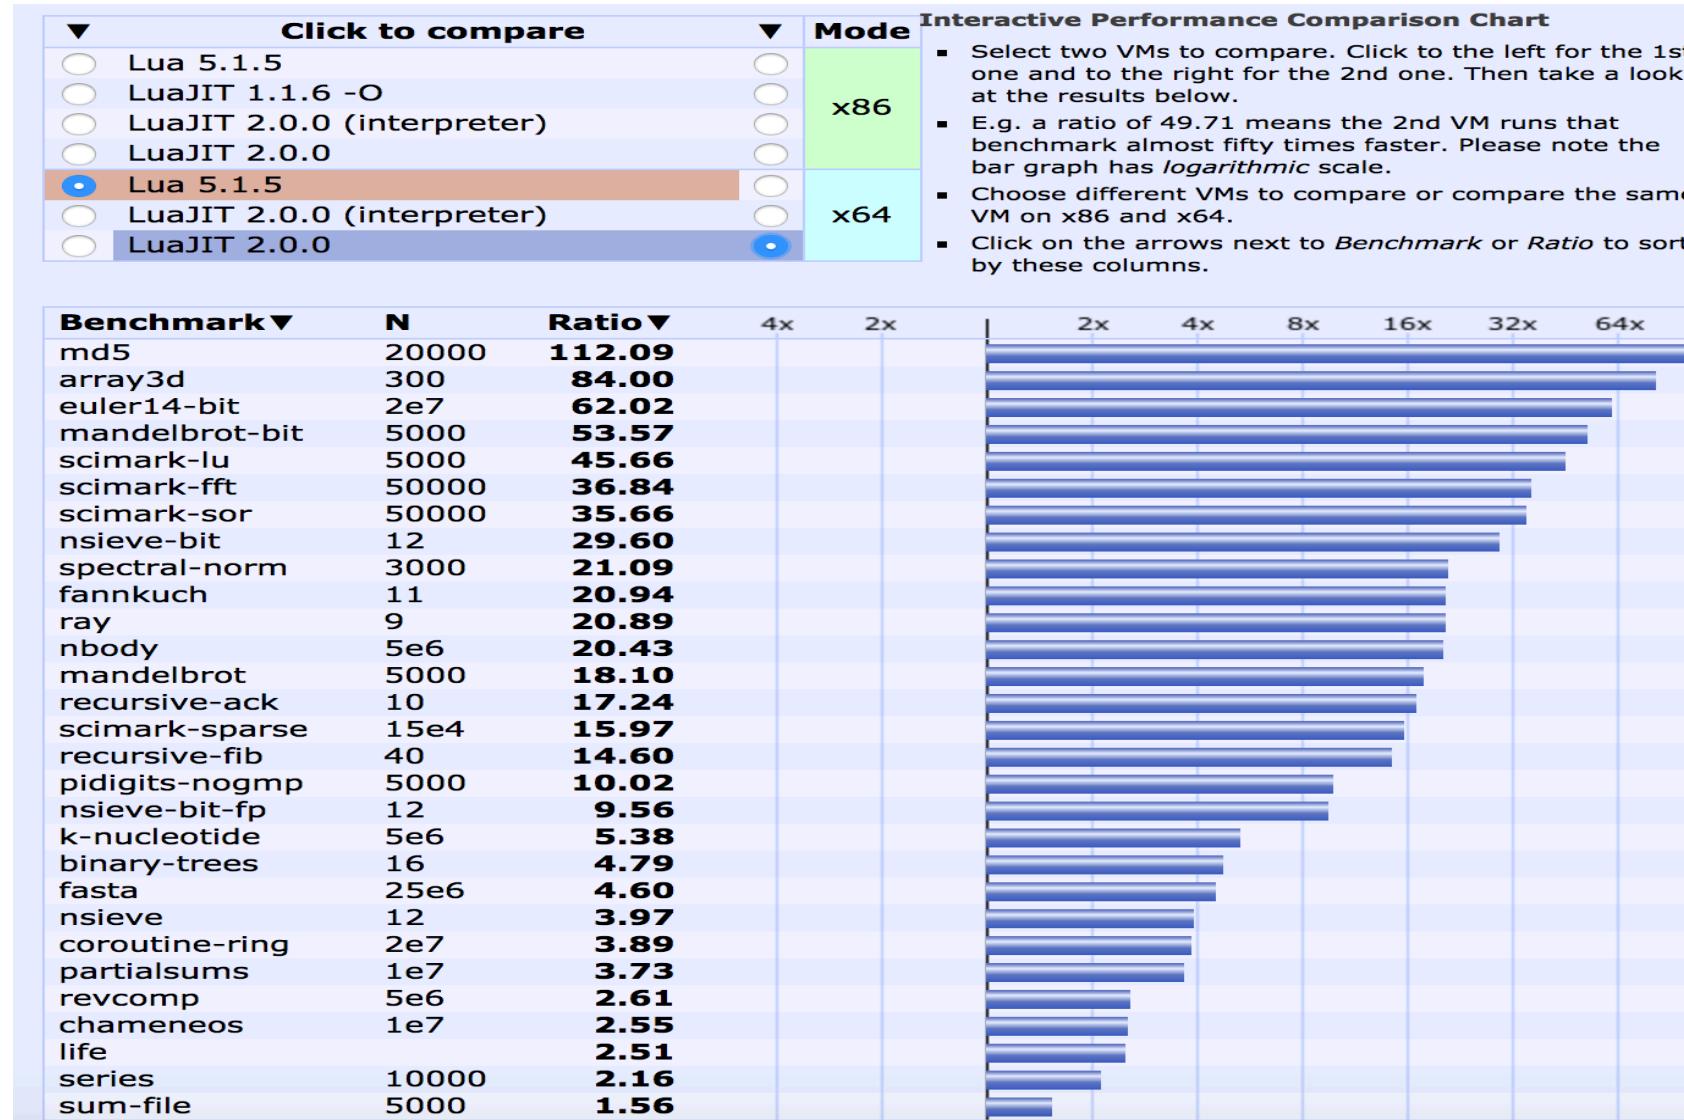
- Lua is a powerful, efficient, lightweight, embeddable scripting language.
- It supports procedural programming, object-oriented programming, functional programming, data-driven programming, and data description Brief description or definition of topic or project
- "Lua" (pronounced LOO-ah) means "Moon" in Portuguese
- Lua is designed, implemented, and maintained by a [team](#) at [PUC-Rio](#), the Pontifical Catholic University of Rio de Janeiro in Brazil
- Lua is [distributed](#) in a small package and builds out-of-the-box in all platforms that have a standard C compiler and supports embedded to IBM Mainframe

Lua/ NGINX (OpenResty) Offers flexibility of Lua with the power of statically compiled C++

- **LuaJITBinding:ZeroMQBinding** (<http://zeromq.org/bindings:lua>)
- **LuaJIT2:**
 - mean throughput: 6,160,911 [msg/s]
 - mean throughput: 1478.619 [Mb/s]
- **C++code:**
 - mean throughput: 6,241,452 [msg/s]
 - mean throughput: 1497.948 [Mb/s]



The Lua JIT (Just In Time) Compiler Ensures that our code runs fast



- **World of Warcraft** : Multiplayer game
- **Adobe Lightroom**: environment for the art and craft of digital photography
- **LEGO Mindstorms NXT**
- Space Shuttle Hazardous Gas Detection System: *ASRC Aerospace, Kennedy Space Center*
- **Barracuda Embedded Web Server**
- **Wireshark** : Network protocol analyzer
- Asterisk: Telecom PBX
- Radvision SIP Stack
- **Redis, RocksDB**(Facebook), Tarantool and many other DBs
- Capital One:
 - Capital One DevExchange API Gateway,
 - Virtual Cards
 - Tokenization Platform

Capital One's Restful API & Architecture Journey

Capital One has been investing heavily in RESTful APIs since 2014

- There was a strong need for a gateway to serve as the single point of entry for all API traffic.
- The Gateway handles
 - Authentication
 - Authorization
 - Rate Limiting
 - API routing
 - Custom Policies

By 2016 we had an opportunity to consolidate a number of legacy gateway products

- Given our GW consolidation & migration strategy, our requirements grew complex



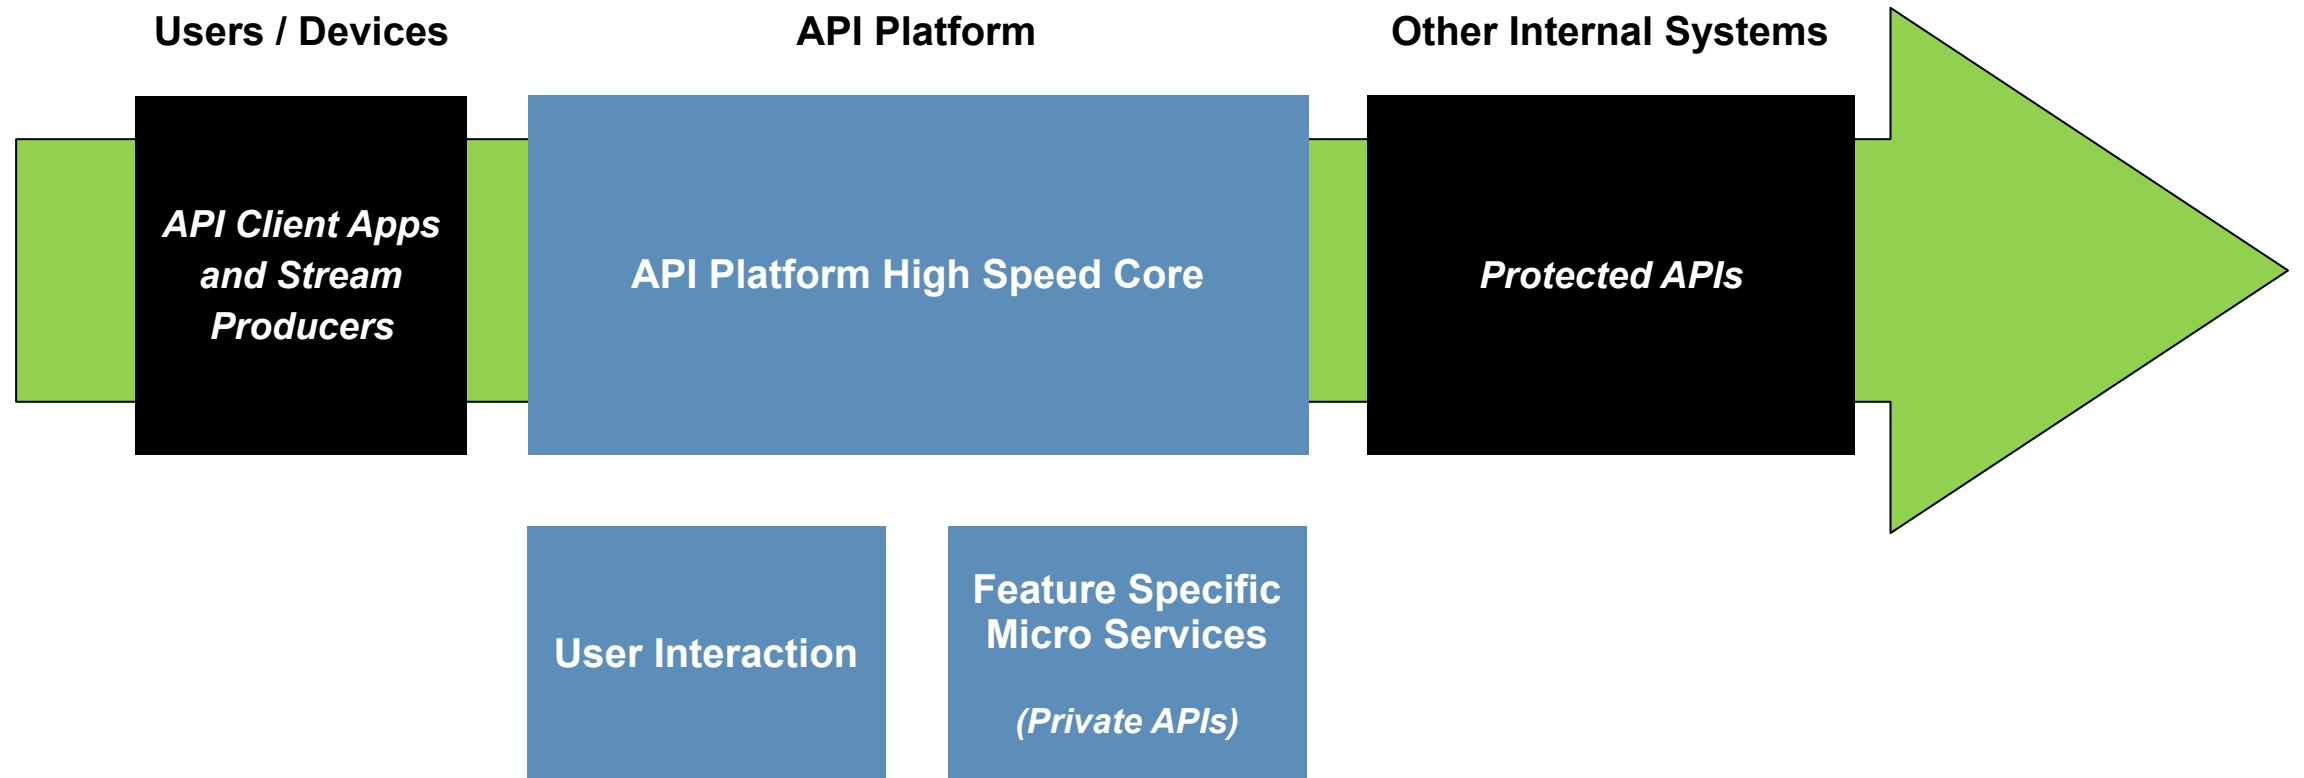
- 12 external options were evaluated
 - 8 options were eliminated based prior to load testing
 - 4 commercial & open source gateways were load tested head to head
- In addition, we evaluated Rohit's Prototype
- We selected our home-grown solution based on features, performance, resiliency and scalability

At first I didn't believe Rohit, so I did my own testing

Experimental API to test out the relative performance

Language (Framework)	Multitasking Model Used	Average Throughput
NodeJS (ExpressJS)	Single Threaded Event Loop	~12K TPS
Java (Spring Boot)	Multi-threaded	~15K TPS
Go (Standard Libraries)	GO Routines	~95K TPS
Lua JIT with NGINX	Single Threaded Event Loop	~97K TPS

We separate features based on the Level of performance required



We defined our Architecture / Design Principles to ensure we can meet our high NFRs

- Leverage ACID transactions only where required and avoid them where possible
- Make systems stateless or leverage Immutable data that is safe to cache indefinitely
- Separate reads from writes
- Partition or Shard Data to meet SLAs
- Micro-batch processing

We Leverage ACID* Transactions Only Where Required and Avoid Them Where Possible

- Ensuring data consistency is hard:
- Data replication and coordination takes time
- Examples Requiring ACID Properties:
 - Issuing Virtual Credit Card Numbers, Issuing New Tokens & Coordinating API changes
- Examples that don't require ACID Properties:
 - Logging, Reading of Immutable Data/Tokens



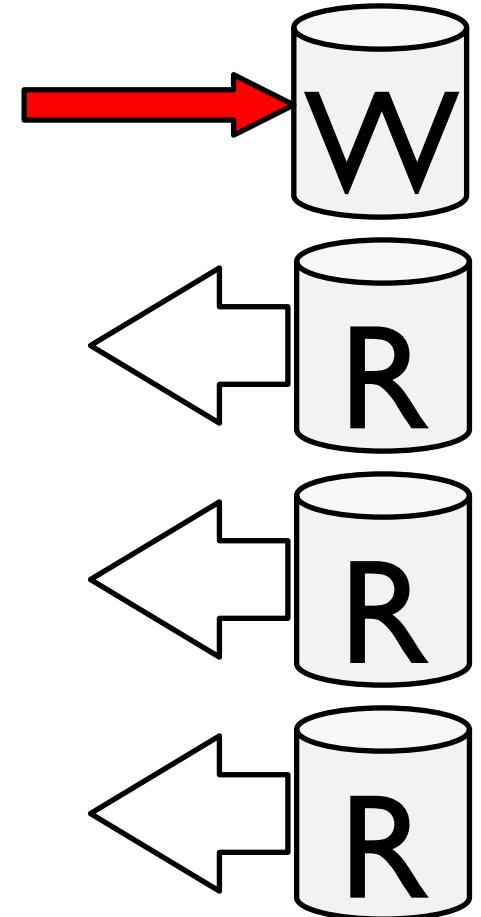
We Make Systems Stateless or Leverage Immutable Data that is Safe to Cache Indefinitely

- **Many API Gateways store a copy of Access Tokens in a database**
- **The Token Lifecycle can be broken into 2 pieces to make it scale better:**
 - DevExchange Gateway Issues Stateless JWE Access Tokens
 - Revoking an Access Token can still be Accomplished with a token blacklist
- **For the Tokenization Use Cases are immutable and can be cached permanently on each server**

We Separate Reads from Writes to Scale

- Separating Reads and Writes can Allow them to be scaled differently without inhibiting the other operation
- For the Tokenization use case:
 - The relationship between Tokens and Original Values are cached on every machine
 - Creating new tokens requires ACID transactions and uses RDS underneath with out of region encrypted read replicas

Workload: 98% Reads / 2%Writes



Partition or Shard Data to meet SLAs

- **Partitioning or sharding data can:**
 - Spread the load
 - Guarantee cache availability
 - Ensure consistent performance
- **Partitioning can be managed manually or provided by the Storage Platform**
- **Tokenization Use Case:**
 - Data is partitioned based on field type (Separate Caches and RDBMs)



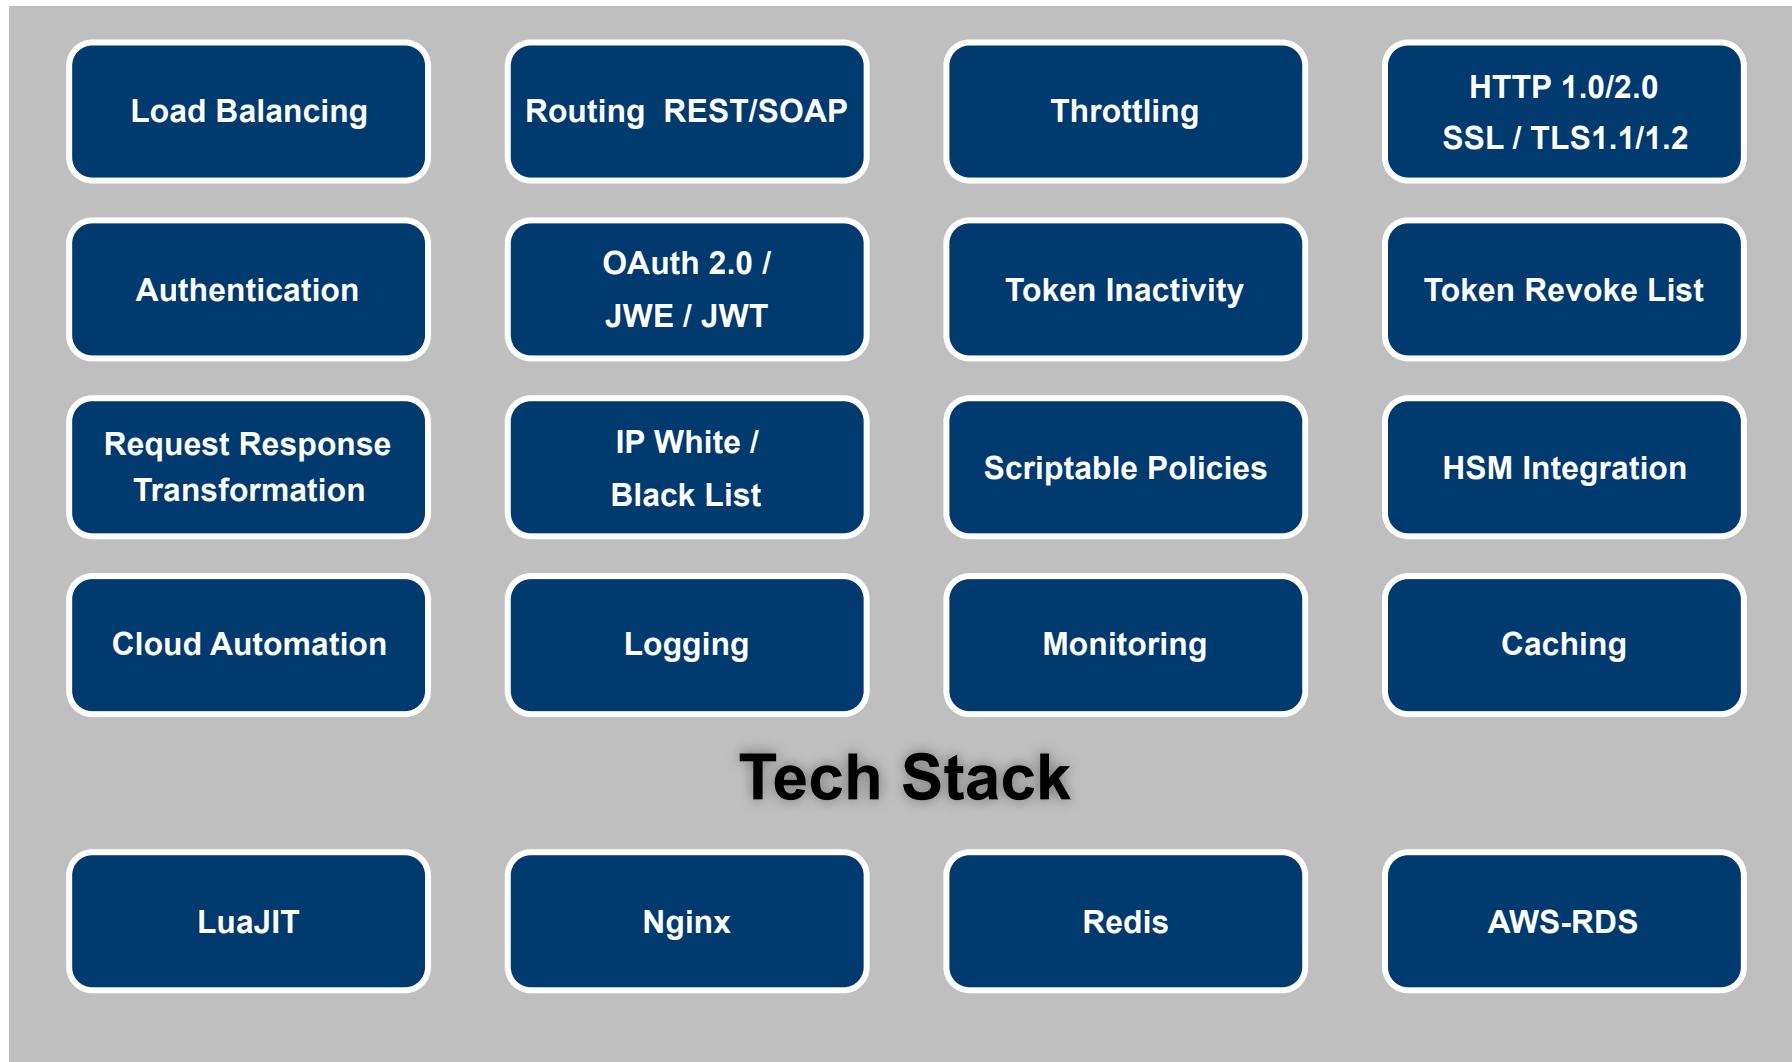
Performance Testing

On the same test bed different load testing tools gave very different results:

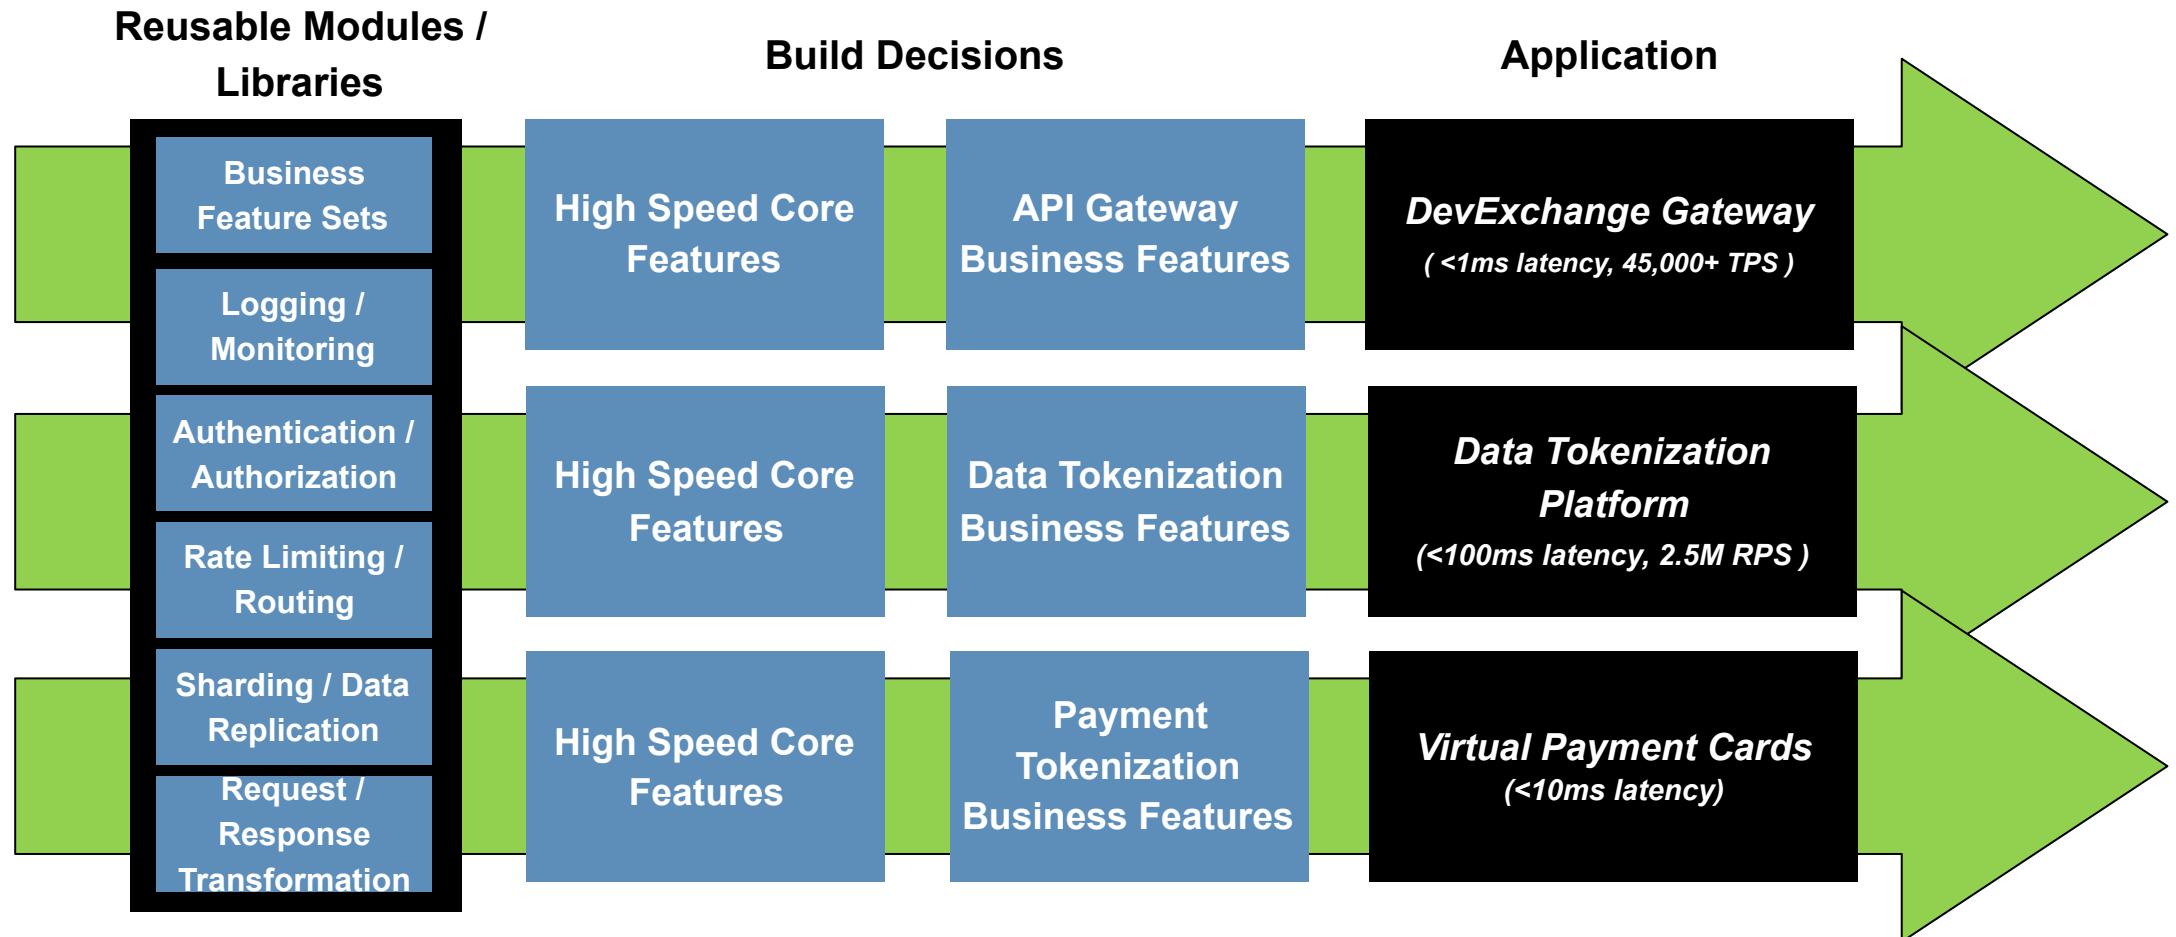
- SOAP UI: 816 TPS
- AB: 7,088 TPS
- JMeter: 24,332 TPS
- WRK: 50,646 TPS

Results

We were able to deliver a full featured API Gateway product to our stakeholders



Based on the success of the API Gateway, we built 2 additional systems on the same stack



Our NGINX stack has enabled us to meet and exceed all of our expectations

- **DevExchange Gateway**

- >2 billion transactions per day
- 45,000 transactions per second (peak)
- < 1ms latency (Average)

- **Data Tokenization Platform**

- 4+ Billion records
- 3+Terabyte of data
- 12 billion operations per day
- 2.5 million operations per second (peak)
- 20 – 40ms latency (Average)

- **Virtual Payment Cards**

- < 2 ms latency (Average)

Thank You